TDB-ACC-NO: NN9504427

DISCLOSURE TITLE: Programmable Thresholds for Reducing Message Latency

PUBLICATION-DATA: IBM Technical Disclosure Bulletin, April 1995, US

VOLUME NUMBER: 38

ISSUE NUMBER: 4

PAGE NUMBER: 427 - 428

PUBLICATION-DATE: April 1, 1995 (19950401)

CROSS REFERENCE: 0018-8689-38-4-427

DISCLOSURE TEXT:

This document contains drawings, formulas, and/or symbols that will not appear on line. Request hardcopy from ITIRC for complete article.

Disclosed is a programmable threshold mechanism for the passing of message packets among nodes of a parallel system at reduced latencies.

The state-of-the-art solutions use store-and-forward methods to pass a message across an inter-connection network. Local FIFOs on each side of the network are used to receive and store the entire message before the message can be forwarded to the next FIFO or node. This introduces increasing latency as the message sizes get larger.

The new feature disclosed replaces the store-and-forward approach with a programmable FIFO threshold which permits the FIFO to begin forwarding a message sooner, when it has partially received a message. A programmable threshold can be used to trigger the FIFO to start forwarding a message when a predetermined amount has been stored to the FIFO. The threshold can be set by the node software to send a message when the FIFO fills above any threshold - 1 word, 2 words, n words, etc. The result is that writing and reading of the FIFO can be overlapped (pipelined) to cut transfer latency of the message to a faction of that of the store-and-forward method. Using the threshold method, a message will commence to be forwarded once the FIFO fills

to
the threshold level and will continue until the message is complete,
even if the FIFO subsequently falls to below the threshold level. In
addition, any message which is too short to cross the threshold upon
being completely written to the FIFO, still triggers the normal
store-and-forward operation and causes the message to be forwarded
immediately.

The Figure shows a typical programmable threshold implementation where the threshold value is written by software into
a
hardware register. The FIFO is a two-ported RAM that can be written
and read simultaneously over separate ports. Each port is controlled
by a pointer, which locates where the next data is to be written or
read. The write pointer is incremented as each data word is written
to the FIFO, and the read pointer is incremented as each data word is
read from the FIFO. A subtractor continuously subtracts the read
pointer from the write pointer to determine how many data words are in
the FIFO at any given time.
The substractor output is compared to
the
threshold register. When the comparator indicates that the substractor
output is compared to the threshold register. When the comparator
indicates that the subtractor output has become greater than the
threshold value, the read enable latch is set and the forwarding
process commences. Only a message that has been completely read will
cause the read enable latch to reset, so that even if the FIFO level
falls below the threshold, the read operation will continue.

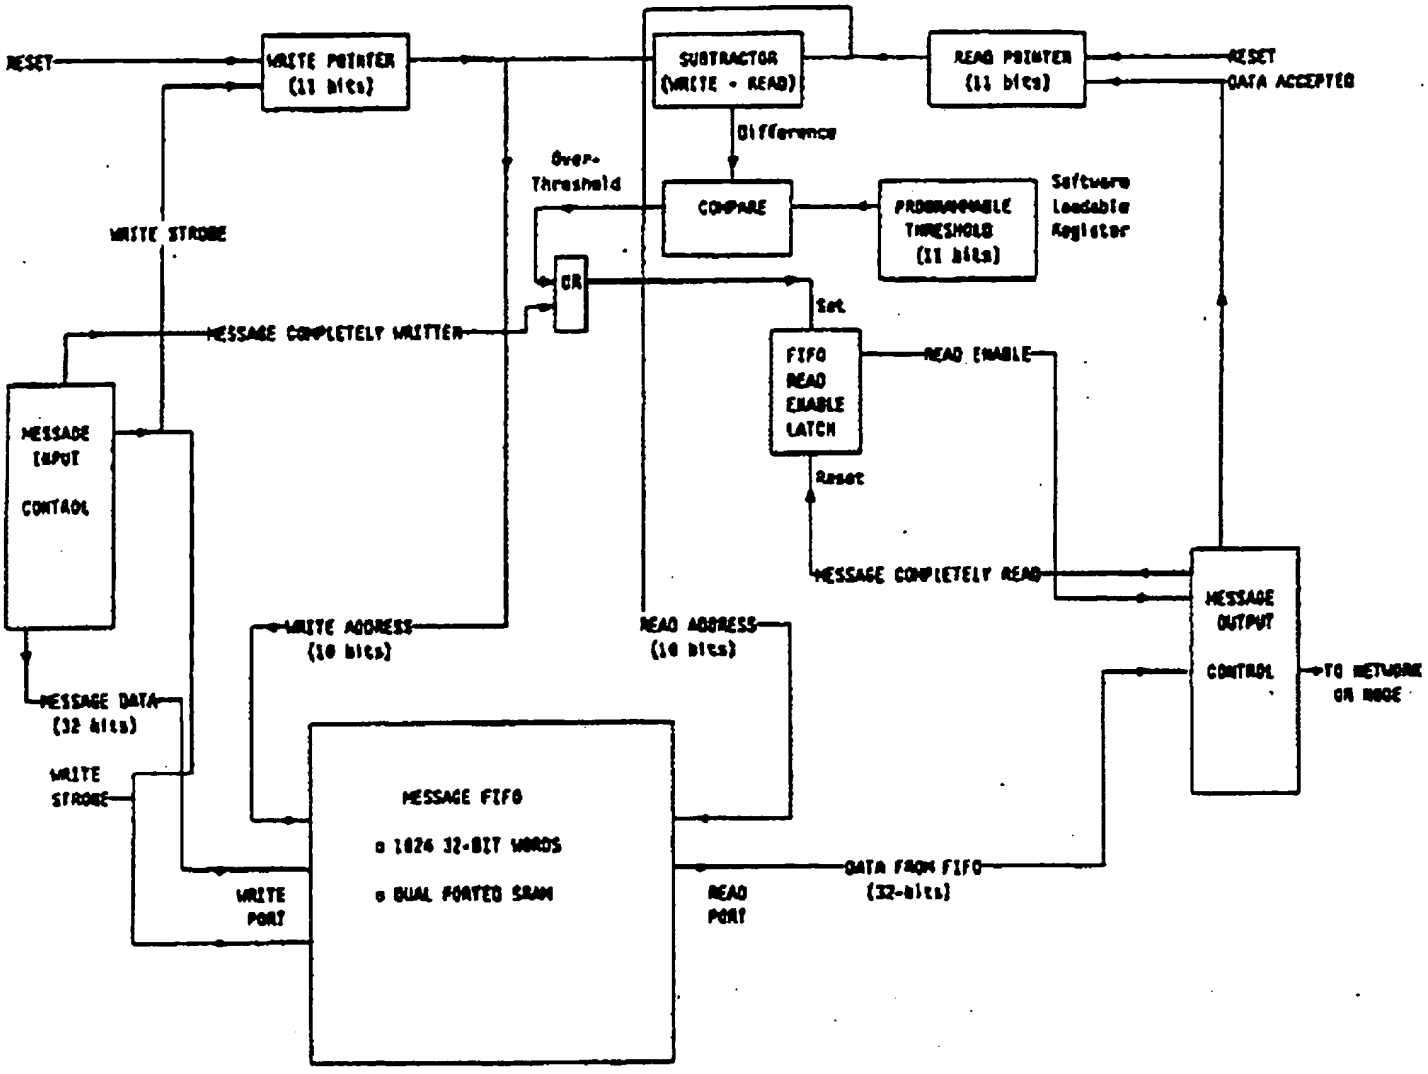
RESET
WRITE POINTER (11 bits)
SUBTRACTOR (WRITE - READ)
READ POINTER (11 bits)
RESET
DATA ACCEPTED
Difference
Over-Threshold
COMPARE
PROGRAMMABLE THRESHOLD (11 bits)
Software Loadable Register
WRITE STROBE
CR
Set
MESSAGE COMPLETELY WRITTEN
FIFO READ ENABLE LATCH
READ ENABLE
MESSAGE INPUT CONTROL
Reset
MESSAGE COMPLETELY READ
MESSAGE OUTPUT CONTROL
TO NETWORK OR NODE
WRITE ADDRESS (10 bits)
READ ADDRESS (10 bits)
MESSAGE DATA (32 bits)
WRITE STROBE
MESSAGE FIFO
• 1024 32-BIT WORDS
• DUAL PORTED SRAM
WRITE PORT
READ PORT
DATA FROM FIFO (32-bits)